

Using HB Family of Protocols for Privacy-Preserving Authentication of RFID Tags in a Population

Tzipora Halevi¹, Nitesh Saxena¹, and Shai Halevi²

¹ Polytechnic Institute of New York University

² IBM T. J. Watson Research Center

Abstract. In this paper, we propose an HB-like protocol for privacy-preserving authentication of RFID tags, whereby a tag can remain anonymous and untraceable to an adversary during the authentication process. Previous proposals of such protocols were based on PRF computations. Our protocol can instead be used on low-cost tags that may be incapable of computing standard PRFs. Moreover, since the underlying computations in HB protocols are very efficient, our protocol also reduces reader load compared to PRF-based protocols.

We suggest a tree-based approach that replaces the PRF-based authentication from prior work with a procedure such as HB+ or HB#. We optimize the tree-traversal stage through usage of a “light version” of the underlying protocol and shared random challenges across all levels of the tree. This provides significant reduction of the communication resources, resulting in a privacy-preserving protocol almost as efficient as the underlying HB+ or HB#.

1 Introduction

Radio Frequency Identification (RFID) technology is increasingly used in many aspects of daily life. Low-cost RFID devices have numerous applications in military, commercial and medical domains. In most of these applications, an RFID reader must authenticate its designated tags to prevent tag forgery and counterfeiting. At the same time, privacy concerns in many of these applications dictate that the tag’s identity is not revealed to an attacker, even while the tag authenticates itself (and even if the attacker pretends to be a reader).

RFID tags are typically very low-cost devices and their computation and storage capabilities are severely constrained. Hence, traditional authentication protocols based on Pseudo-Random Functions (PRFs) may not be applicable on such tags. To this end, Juels and Weis [9] suggested using the HB protocol (of Hopper and Blum [8]) for tag authentication. Several enhancements such as HB+ [9] and HB# [6] were proposed later, and are currently the only viable solutions for very low-cost authentication.

Further complicating RFID authentication is the fact that a single reader must often handle a very large tag population, possibly in the millions or even billions. Clearly, to use the appropriate secret keys, the reader must learn the identity of a tag that is trying to authenticate itself. Although the tag can identify itself by sending a pre-shared unique identifier, such identifiers can easily be used by an eavesdropping adversary to compromise the privacy and anonymity of the tags (and thus also of their bearers). Therefore, a privacy-preserving solution is needed to allow the reader to identify the tag.

One possible approach to achieve *privacy-preserving authentication* is for the reader to perform an exhaustive search, by trying to match the result of the tag being authenticated

against all possible tags. However, this requires a large $O(N)$ amount of computation for a population of N tags. To address this scalability problem, Molnar et al. [14] proposed using a “tree of keys”, assigning to each tag the keys corresponding to a root-to-leaf path in the tree, and using these keys as seeds for a PRF; the PRF-based response of a tag is then linked to this path, and this link is used to identify the tag. This solution reduces the reader computational complexity of $O(\log(N))$. Also, Cheon et al. [4] proposed a 2D-mesh scheme that enables the reader to identify the tag with complexity of $O(\sqrt{N} \log N)$: The idea is to create two sets of PRF keys, each with \sqrt{N} keys, and give every tag one key from each of the two sets, such that the combination of these two keys uniquely determined that tag. In [2], Burmester et al. also utilized PRF functions for an anonymous RFID authentication scheme with constant lookup. In this approach, the reader keeps for every tag a constant number of pseudonyms. At the end of each successful authentication session the tag generates a new pseudonym, which the reader then uses to identify the tag from its tag lookup pseudonyms table.

All the above proposals are based on protocols that require the tag to compute PRF operations and therefore not applicable to very low-cost tags. To this end, we consider the HB family of protocols for privacy-preserving authentication. Adapting HB-like protocols to this setting takes some care, however. For example, HB protocols are typically designed with very tight parameters, which yield a non-negligible false accept rate (FAR), i.e., the probability of an illegitimate tag successfully authenticating to the reader. In an exhaustive search, the overall false accept rate grows roughly by a factor of N . For example, an instance of the HB+ protocol with 80 rounds and noise probability of 0.25, has an FAR of 4×10^{-6} . When using this instance, with an exhaustive search over a population of a million tags, the overall FAR would be close to one, which is clearly unacceptable for any practical use.

1.1 Our Contributions

In this work we develop a tree-based protocol for privacy-preserving authentication. Specifically, we study the use of HB-like protocols (such as HB+ or HB#). Our tree-based protocol has two logical stages. First there is a *tree traversal stage*, in which the reader tries to identify the most likely tag to authenticate, and then there is an *authentication stage*, in which the identity of that “most likely tag” is verified. Unlike exhaustive search, the false-accept rate of this protocol does not grow with the number of tags in the population, since only one tag is authenticated in the last stage. We also show that the false-reject rate grows very slowly, despite the presence of a large number of tags. We propose several optimizations over naive use of HB+/HB# in a tree-based scheme. These optimizations bring the communication of the tree-based protocol down to not much more than the underlying HB+/HB# authentication protocol, and also improve the computation time.

Our work is based on the observation that the function $f_x(A) = Ax + \text{noise}$ for a random matrix A behaves like a (randomized) pseudo-random function (which was proven in [1, Lemma 3.2] under the assumption that learning parity with noise is hard). Therefore we replace the standard PRFs during the tree-traversal stage with this “HB-like” function. We observe that this function only needs tag-generated random challenges during the tree-traversal stage. (It is only in the authentication stage that a reader-generated challenge is needed.) Moreover, there is no need to generate a new challenge

in every level, it is perfectly safe to use the same random challenge in all the levels of the tree, and we can use significantly smaller parameters during the tree-traversal stage. The combination of these observations reduces computation and communication by a factor of $2\times$ to $4\times$ (In fact our tree protocol is almost as efficient as the underlying HB+ or HB#). We present analytical and simulation results comparing our method with prior proposals in terms of computation, communication and memory overheads.

We note that just like all other HB-type protocols, ours is also vulnerable to man-in-the-middle attacks [7, 16]. Also, just like other tree-based protocols, its privacy guarantees degrade somewhat when the secret keys of many tags are exposed to the adversary.

2 Background: HB-type Protocols

In the original HB protocol by Hopper and Blum [8], the reader and the tag share a secret \mathbf{x} which is a k -bit vector (e.g. $k = 256, 512$, etc.). The protocol consists of many rounds. In each round the reader sends a random challenge vector \mathbf{a} and the tag replies with one bit $z = \mathbf{a} \cdot \mathbf{x} \oplus \nu$, where ν is 1 with probability ε (for some parameter $\varepsilon < 0.5$). After r such rounds, the tag is accepted if z matches $z' = \mathbf{a} \cdot \mathbf{x}$ at least $r - \tau$ times, where τ is some threshold parameter.

The HB protocol is vulnerable to attacks by an active adversary, who can send to the tag the same challenge \mathbf{a} many times, thereby eliminating the noise and recovering the secret \mathbf{x} . To counter this attack, the HB+ protocol by Juels and Weis [9] added a second secret \mathbf{y} that is shared by the reader and the tag. In each round of the protocol, a random challenge \mathbf{b} is generated by the tag, a second challenge \mathbf{a} is generated by the reader, and the tag sends to the reader $z = \mathbf{a} \cdot \mathbf{x} \oplus \mathbf{b} \cdot \mathbf{y} \oplus \nu$. As before, the tag is accepted after r rounds if z matches $z' = \mathbf{a} \cdot \mathbf{x} \oplus \mathbf{b} \cdot \mathbf{y}$ at least $r - \tau$ times.

It was later shown in [10, 11] that all the rounds can be carried out in parallel: The tag chooses a matrix B , the reader responds with a matrix A , the tag chooses a noise vector ν and replies $\mathbf{z} = A \cdot \mathbf{x} \oplus B \cdot \mathbf{y} \oplus \nu$, and the reader checked that \mathbf{z} is close enough to $\mathbf{z}' = A \cdot \mathbf{x} \oplus B \cdot \mathbf{y}$. This protocol can withstand an active attacker that interacts with the tags, but it is still vulnerable to man-in-the-middle attacks [7, 16].

The HB# protocol [6] is a variant where the “types” of the secrets and the challenges are swapped: Namely, the tag and reader share secrets X and Y which are $r \times k$ matrices, the challenges are r -vectors \mathbf{a} and \mathbf{b} , and the reply that the tag computes is $\mathbf{z} = \mathbf{a} \cdot X \oplus \mathbf{b} \cdot Y \oplus \nu$. To alleviate the additional storage requirement, HB# proposes using Toeplitz matrices for X, Y . Another claimed benefit for swapping the types is that HB# can resist some limited form of man-in-the-middle attacks, however, not all [16].

2.1 Security

The security of the HB series of protocols relies on the computational hardness of the problem of *Learning from Parity with Noise* (LPN) [9]. Roughly, the LPN problem is to find a secret k -vector \mathbf{x} given a random binary $r \times k$ matrix A and a vector $\mathbf{z} = A \cdot \mathbf{x} \oplus \nu$, where ν is a noise vector in which each entry is set to 1 with probability ε (for a parameter $0 < \varepsilon < \frac{1}{2}$).

The hardness of this problem (and thus the security level of the protocols) depends on the choice of the parameters ε and k . The best algorithm for learning parity with noise is due to Blum, Kalai and Wasserman [3], with some later optimizations, notably

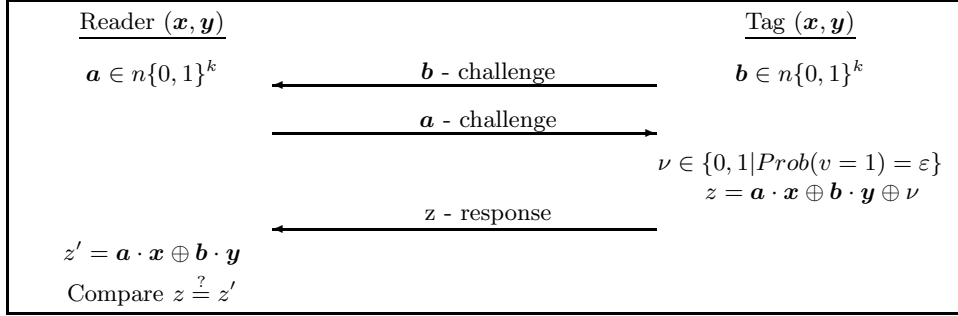


Fig. 1. One round of the HB+ Protocol

by Leveil and Fouque [12]. Based on the performance of the BKW algorithm, Juels and Weis suggested in [9] that setting $k = 224$ and $\varepsilon \in [\frac{1}{8}, \frac{1}{4}]$ yields an 80-bit security level. In [12], Leveil and Fouque claimed, however, that these parameters yield only a 55-bit security level, and that with $\varepsilon = \frac{1}{4}$ one needs to use $k = 512$. (They also noted that only one of the secrets \mathbf{x}, \mathbf{y} in HB+ needs to be this long, while the other can safely be only 80-bit long.)

The other parameters of interest in HB-type protocols are the number of rounds³ and the acceptance threshold. Juels and Weis suggested in [9] to use number of rounds $r \in [40, 80]$ and set the threshold at $\tau = \varepsilon r$. This gives a relatively high false reject rate (FRR) of 0.44, and when used with error probability $\varepsilon \in [\frac{1}{8}, \frac{1}{4}]$ it gives false-accept rates (FAR) between 10^{-3} and 10^{-9} . (The HB# proposal in [6] suggested other parameters with a much lower FRR and FAR, e.g., $\text{FAR} \leq 2^{-80}$ and $\text{FRR} \approx 2^{-45}$.)

3 Private Authentication with an HB-like protocol

Recall that in our setting, we have a system with many tags under one administrative domain. Similarly to [14], we arrange all the tags in a tree structure and use this structure to find the right tag to authenticate. However, we use an HB-like procedure to implement both the tree search and the tag authentication as opposed to the PRF-based solution in [14].

A naive approach would simply perform the underlying authentication protocol at every level of the tree. Namely, at every level one can run HB+ (or HB#) and the reader then performs an exhaustive search over the children of the current node, selecting the one for which the authentication was successful.⁴ A closer inspection reveals that this naive approach can be significantly improved:

- First, there is no reason to use challenges from the reader while traversing the tree. Indeed, the challenge from the reader is needed for secure authentication, but *not* for the identification of “the right tag to authenticate”.

³ The name “number of rounds” refers to a sequential protocol where a one-bit z is returned in every round. For the parallel version, this parameter is the reply-size.

⁴ In the rest of this paper, we concentrate only on the case of HB+. The treatment of HB# is completely analogous.

We therefore suggest to logically split the protocol into a *tree-traversal stage*, where the reader finds the right tag to authenticate, and an *authentication stage*, where the identity of that tag is verified. In the authentication stage, we just use the underlying HB+ protocol unchanged, but in the tree-traversal stage we only use random challenges from the tag. This optimization cuts the computation time and storage requirements for both the tag and the reader.

(We note, however, that in the asymmetric setting suggested in [12], where one secret is longer than the other, we need to store a long secret in every internal node in the tree. For example, using the 512/80 asymmetric setting, this optimization only saves 80 of the combined $512 + 80 = 592$ bits, which is $\approx 13\%$.)

- More importantly, we observe that we can share the same challenges across all the levels. Namely, instead of the tag sending a different challenge for every level, we send only one challenge and use it at all the levels. This means that we only have communication of two challenges throughout the tree-based protocol (just as in the underlying HB+), and the only additional communication is an r -bit response vector for every level.

In more details, the reader in our system maintains a tree structure with the tags at the leaves, and each internal node n in the tree is associated with a secret key \mathbf{y}_n , which is a binary vector of length k_y . A tag t , associated with a leaf in the tree, is given all the secret keys that are associated with nodes on the path from the root to that leaf. In addition the tag is also given two unique secret keys \mathbf{x}_t and \mathbf{y}_t (of length k_x and k_y , respectively) that are used for authentication. The notations that we use in the description below are summarized in Table 1.

Table 1. Notation table

d	depth of the tree
β	branching factor of the tree
k_x, k_y	length of the secrets in the system ($k_x = 80, k_y \in [224, 512]$)
r	size of the “response” messages (usually $r \in [80, 128]$)
ε	noise level ($\varepsilon = 0.25$)
τ	acceptance threshold (usually $\tau \in [20, 40]$)

3.1 System Setup

The system is setup with some upper bound N on the number of tags that it can support, and the parameters r, d, β are derived from this upper bound. (See Section 5.6 for a discussion of how to determine these parameters.) Once determined, the parameters d, β define a tree with $\beta^d \geq N$ leaves, which will be associated with the tags in the system. Also, a “master secret” is chosen, which will be used to determine all the other secrets in the system as described next.

3.2 Tag Registration

When a tag is registered into the system, it is associated with a random available leaf in the tree. This can be done either by the system administrator remembering the positions

of all the tags in the tree and choosing at random one leaf that is not yet used, or by choosing a random permutation π over the domain $[1, \beta^d]$ and assigning the i 'th tag in the system to leaf $\pi(i)$.⁵

Let n_0, n_1, \dots, n_d be the path in the tree from the root (denoted n_0) to the leaf that is associated with the new tag (denoted n_d). For each node n_i (except the root), the tag is given the k_y -bit secret \mathbf{y}_{n_i} that is associated with that node. (This secret can be derived from the master secret, say by setting $\mathbf{y}_{n_i} = \text{PRF}_{ms}(n_i)$ for some pseudo-random function PRF that is keyed by the master secret ms .) The tag also gets two additional keys \mathbf{x}_t and \mathbf{y}_t of size k (that can similarly be derived from the master secret). After registration, the tag will hold the keys $\mathbf{y}_{n_1}, \dots, \mathbf{y}_{n_d}, \mathbf{x}_t, \mathbf{y}_t$.

3.3 Tree-Traversal Stage

At the start of the authentication process, the tag needs to be identified so the correct keys can be used for the authentication. To this end, the tag chooses a $r \times k_y$ random challenge matrix B , and noise vectors ν_i for every level i in the tree. The tag computes $\mathbf{z}_i = B \cdot \mathbf{y}_{n_i} \oplus \nu_i$ and sends to the reader the matrix B and all the \mathbf{z}_i 's.

Starting from the root, the reader then goes down node by node, using the \mathbf{z}_i 's to decide what child of the current node it needs to use next. Specifically, for every child c of the root it computes $\mathbf{z}_c = B \cdot \mathbf{y}_c$, where \mathbf{y}_c is the secret associated with that child. Then, the reader descends into the child c for which \mathbf{z}_c is closest to the response vector \mathbf{z}_1 received from the tag. Similarly, after descending into an internal node n_i at level i , the reader computes $\mathbf{z}_c = B \cdot \mathbf{y}_c$ for every child c of n_i , and descends into the child c for which \mathbf{z}_c is closest to \mathbf{z}_i that was received from the tag. (Throughout this process, if two children are equally close to \mathbf{z}_i , one is chosen arbitrarily as the next node).

After using all the \mathbf{z}_i 's, the reader arrives at one of the leaves of the tree, and then it uses the secrets \mathbf{x}, \mathbf{y} that are associated with the tag at that leaf in the authentication stage, as described next.

3.4 Authentication Stage

At the end of the Tree-traversal stage, the reader arrives at a leaf that it considers to be the most likely to be the one associated with the correct tag. At this stage, we need to run the authentication protocol to confirm that the tag is valid. Here we just use the parallel HB+ protocol. Specifically, the reader sends a random challenge matrix A , the tag chooses a noise vector ν , and reply with $\mathbf{z} = A \cdot \mathbf{x}_t \oplus B \cdot \mathbf{y}_t \oplus \nu$ (using the two last secret keys $\mathbf{x}_t, \mathbf{y}_t$ that it holds). The reader recovers the keys \mathbf{x}, \mathbf{y} that are associated with the tag in this most likely leaf, computes $\mathbf{z}' = A \cdot \mathbf{x} \oplus B \cdot \mathbf{y}$, and checks that the Hamming distance between \mathbf{z} and \mathbf{z}' is below the threshold τ . It accepts the tag if \mathbf{z} and \mathbf{z}' are close enough, and rejects it otherwise.

4 Security Analysis

Recall that we have two security goals in our model, namely privacy and authentication. Our attack model allows the adversary to eavesdrop on the communication between

⁵ An efficient method for implementing such a random permutation was recently proposed by Morris et al. [15].

tags and the reader, and also to communicate directly with the tag and the reader, but NOT to modify messages that are sent between them. In other words, we consider an active adversary, but explicitly disregard man-in-the-middle attacks (since all HB-type protocols are insecure against them [7, 16]).

4.1 Authentication

Our model for authentication is essentially the usual DET model [9, 10, 6], where in our case we have one target tag that the attacker is trying to impersonate, and all the other tags in the system are considered to be adversarial. In more details, the attacker first gets the secrets of all the tags in the system but one, then it can interact with the remaining tag for q times, and finally it interacts with the reader and tries to impersonate that remaining tag.

Since we just run the original HB+ protocol at the authentication stage, then trivially our protocol is as resilient against impersonation as HB+. (Formally, there is a trivial tight reduction from breaking our protocol to breaking the original HB+.)

4.2 Privacy

Our privacy model meant to capture an attacker who can interact with many tags, and who tries to link different authentication sessions. For example, consider an attacker that roams through a library talking to tags and recording the books that these tags are embedded in. Later the attacker detects some tags in the book-bag of a random person on the street, and it tries to recognize these tags (thereby recognizing the books that this person checked out).

A (somewhat simplistic) formalization of this concern has the adversary talking to two tags in the system, which are chosen as follows: First the adversary chooses two arbitrary distinct tags in the system, and then a fair coin is tossed. If it comes out **head** then the two tags are the ones chosen by the adversary, and if it comes out **tail** then the two tags that the adversary talks to are actually just the first tag that it chose. The adversary now interacts with the “two tags” for a total of q sessions, at the end of which it needs to guess the outcome of the coin toss.

We note that this is not the only possible definition for privacy in this setting (and maybe also not the most natural one). For example, this definition does not consider the possibility of the adversary compromising some tags and learning their secrets. Still, we think that this model is a reasonably meaningful one.

The security of our protocol in this model follows easily from the fact that the function $f_{\mathbf{y}}(B) = B \cdot \mathbf{y} + \text{noise}$ is a pseudorandom function when B is random, and this fact was proven in [1, Lemma 3.2] to follow from the hardness of LPN. Indeed, it is easy to see that if we replace this function in our protocol with a truly random function then the adversary would have zero advantage in guessing the outcome of the coin-toss in the experiment from above. (The fact that we can share the same challenge matrix B for all levels of the tree follow from an easy hybrid argument.)

We finally note that our protocol, like all tree-based protocols, is somewhat vulnerable to tag-compromise. If an attacker manages to get the secrets of a specific tag, it can use them to recognize the tree-traversal replies of neighbor tags in the tree (and in particular it can distinguish neighbors of the compromised tags from non-neighbors).

Devising mechanisms to mitigate the effect of tag compromise is an interesting open problem.

5 Parameters and Performance

Below we analyze the false reject and false accept rates for a fixed set of parameters β, d, r , and then use this analysis to set these parameters (as a function of the total number of tags N).

5.1 False-Accept Rate

We begin by observing that the false-accept rate of our protocol is exactly the same as that of the underlying HB+ that is used in the authentication stage. To see why, observe that the tree-traversal stage of the protocol always results in the reader identifying one leaf of the tree as the most likely to correspond to the right tag, and then the underlying HB+ is run against the keys of the tag in this most likely leaf.

This is in sharp contrast to the situation for the trivial linear search routine that was discussed in the introduction. In that procedure, the reader would try to check the authentication against the keys of all the tags in the system, so the false-accept rate would roughly be multiplied by the number of tags. This does not happen in our algorithm, since the reader only tries to authenticate against the keys of just one tag.

In other words, even giving the adversary “for free” the ability to steer the tree-traversal stage to any tag of its choice, it would still have to be able to authenticate a tag without knowing its keys in order to induce a false accept event.

5.2 False-Reject Rate

We note, however, that the false reject rate of our algorithm will be higher than the standard HB+. This is because the tree-traversal stage may identify wrong leaf, in which case the authentication stage will almost surely reject. We therefore need to analyze the probability that the tree-traversal stage identifies the wrong leaf.

Taking a false branch, the binary case. We start by considering a binary tree (i.e., $\beta = 2$) and analyzing the probability that the wrong branch is chosen at one step of the tree-traversal stage. For the binary tree, if the algorithm is currently in a node at level $\ell - 1$ which is on the path to the true tag, then the likelihood of choosing the wrong child for level ℓ is exactly the probability that the non-matching key of that “false child” generates a vector \mathbf{z}_f that is closer to the response vector \mathbf{z}_ℓ than the matching key of the “true child”.

Below we denote by Hwt the Hamming weight of a binary vector. We denote the “true child” of the current node by t (i.e., the child on the path to the true tag that is trying to authenticate itself), and denote the other child of that node by f (for a “false child”). With the keys of these children denoted \mathbf{y}_t and \mathbf{y}_f , respectively, and the tag sending a challenge matrix B , recall that the reader computes $\mathbf{z}_f = B \cdot \mathbf{y}_f$ and $\mathbf{z}_t = B \cdot \mathbf{y}_t$, and compares these two vectors against the vector \mathbf{z}_ℓ that was sent by the tag.

The probability of descending into the wrong child f is bounded by the probability that \mathbf{z}_f is closer to \mathbf{z}_ℓ than \mathbf{z}_t , namely $\Pr[\text{Hwt}(\mathbf{z}_f \oplus \mathbf{z}_\ell) < \text{Hwt}(\mathbf{z}_t \oplus \mathbf{z}_\ell)]$. We denote by $P_t(i)$ the probability that $\mathbf{z}_t, \mathbf{z}_\ell$ differ by exactly i positions, and similarly by $P_f(i)$ we denote the probability that $\mathbf{z}_f, \mathbf{z}_\ell$ differ by exactly i positions. Namely,

$$P_t(i) \stackrel{\text{def}}{=} \Pr[\text{Hwt}(\mathbf{z}_t \oplus \mathbf{z}_\ell) = i] = \binom{r}{i} \varepsilon^i (1 - \varepsilon)^{r-i} \quad (\text{where we use } \varepsilon = 1/4)$$

$$P_f(i) \stackrel{\text{def}}{=} \Pr[\text{Hwt}(\mathbf{z}_f \oplus \mathbf{z}_\ell) = i] = \binom{r}{i} / 2^r$$

Then we bound the probability of taking a false branch by

$$\Pr[\text{Hwt}(\mathbf{z}_f \oplus \mathbf{z}_\ell) < \text{Hwt}(\mathbf{z}_t \oplus \mathbf{z}_\ell)] = \sum_{i=1}^r P_t(i) \cdot \left(\sum_{j=0}^{i-1} P_f(j) \right) \quad (1)$$

We can approximate the above expression as follows: Let Δ_{ft} be the difference between the Hamming weight of $\mathbf{z}_t \oplus \mathbf{z}_\ell$ and $\mathbf{z}_f \oplus \mathbf{z}_\ell$,

$$\Delta_{ft} \stackrel{\text{def}}{=} \text{Hwt}(\mathbf{z}_f \oplus \mathbf{z}_\ell) - \text{Hwt}(\mathbf{z}_t \oplus \mathbf{z}_\ell)$$

The random variable Δ_{ft} is the sum of r independent and identically distribution random variables, one for every position in the response vector (and each a random variable over $\{-1, 0, 1\}$). The probability of taking a false branch is then bounded by $\Pr[\Delta_{ft} < 0]$, and by the law of large numbers we can approximate Δ_{ft} by a Normal random variable with the same mean and variance. Recalling that $\text{Hwt}(\mathbf{z}_t \oplus \mathbf{z}_\ell)$ has mean $\mu_t = \frac{r}{4}$ and variance $\sigma_t^2 = \frac{3r}{16}$ and that $\text{Hwt}(\mathbf{z}_f \oplus \mathbf{z}_\ell)$ has mean $\mu_f = \frac{r}{2}$ and variance $\sigma_f^2 = \frac{r}{4}$, we get that Δ_{ft} has mean $\mu = \mu_f - \mu_t = \frac{r}{4}$ and variance $\sigma^2 = \sigma_t^2 + \sigma_f^2 = \frac{7r}{16}$. Hence we have

$$\Pr[\text{false branch}] \leq \Pr[\Delta_{ft} < 0] \approx \text{erfc} \left(\frac{r/4}{\sqrt{7r/8}} \right) = \text{erfc} \left(\sqrt{r/14} \right)$$

For example, if we choose $r = 80$, then we estimate the probability of choosing a false branch as $\Pr[\text{false branch}] \approx 6.97 \cdot 10^{-4}$, while a choice of $r = 144$ gives $\Pr[\text{false branch}] \approx 5.38 \cdot 10^{-6}$.

Taking a false branch, the general case. For a larger branching factor ($\beta > 2$), we can still use an equation similar to Equation (1) for the probability of choosing a false branch, but estimating it becomes harder. Specifically, we replace the probability that one false tag has less errors than the true tag by the probability that at least one of the false tags has less errors:

$$\Pr[\text{false branch}](\beta) = \sum_{i=1}^r P_t(i) \cdot \left[1 - \left(1 - \sum_{j=0}^{i-1} P_f(j) \right)^\beta \right] \quad (2)$$

However, estimating the last expression is harder, since we need to bound the probability that the minimum of several random variables is below zero (and moreover these random variables are highly correlated).

In lieu of an analytical bound, we therefore evaluated the expression from Equation (2) explicitly, and plotted the results in Figure 2. Specifically, fixing the target false-branch probability to some constant (either 0.1 or 0.01), we plot for every branching factor from $\beta = 2$ to $\beta = 10^4$ the smallest response-size r for which the false-branch probability is below that target. (For example, with branching factor of $\beta = 1000$ we need $r \approx 80$ to get false-branch probability of 0.1.) As expected, for a constant probability of false branch, the response-size r grows linearly with the log of the branching factor.




Fig. 2. Branching factor (β) vs. response-length (r), for false-branch probabilities of 0.1 and 0.01.

The overall false reject rate. Recall that a valid tag can be falsely rejected either due to the algorithm choosing a false branch at some point during the tree-traversal stages, or due to a false reject in the authentication stage. The probability of failing the authentication stage is equal to the probability that the error vector has more than τ ones, namely the false-reject rate of the authentication stage is

$$FRR(\text{auth}) \stackrel{\text{def}}{=} \sum_{i=\tau+1}^r \varepsilon^i (1 - \varepsilon)^{r-i} \binom{r}{i} \quad (3)$$

and the combined false-reject rate of the whole protocol is

$$FRR(\text{Tree-HB+}) \approx d \cdot \Pr[\text{false branch}] + FRR(\text{auth}) \quad (4)$$

As mentioned in [10], the false reject rate can be reduced if the tag checks the noise vector ν and only use it if it has at most τ one-bits.

5.3 Computation, Storage and Communication

The computation of the reader during the tree-traversal stage consists of d levels, wherein each level the reader computes β vectors \mathbf{z}_c (one for every child c of the current node)

and compares them to the response-vector \mathbf{z}_i that the tag sent. Then, the computation during the authentication stage consists of a single execution of parallel-HB+. Recalling that every vector \mathbf{z}_c during the tree-traversal stage takes one matrix-vector product to compute and the authentication stage takes two more matrix-vector products (one of which with a smaller $k_x \times r$ matrix), we get that the total computation on the reader side is between $d\beta + 1$ and $d\beta + 2$ matrix-vector multiplications.

On the tag side, the tag computes one matrix-vector product for every level of the tree in the tree-traversal stage, and two more in the authentication stage (again, one of which with a smaller matrix), for a total of between $d + 1$ and $d + 2$ products.

As for communication, the tag sends the challenge matrix B for parallel-HB+ and the response vectors \mathbf{z}_i , it receives the challenge matrix A and then sends \mathbf{z} . Hence the total communication (in both directions) is $r(k_x + k_y + d)$.⁶ Finally, the storage requirement on the tag is $k_y(d + 1) + k_x$ bits.

An important optimization for our protocol is to choose different values for the response-length r in the two stages, so as to get $d \cdot \Pr[\text{false-branch}] \approx FRR(\text{auth})$. If we denote by r_{tr} the response-length in the tree-traversal stage and by r the response-length in the authentication stage, we get total communication of $r(k_x + k_y) + r_{\text{tr}}d$. The effect on computation is even larger: since we typically have $r_{\text{tr}} < r/2$, the total computation is decreased by about a factor of two.

5.4 Iterating the Protocol

We recall that an easy (and cheap) way of reducing the false-reject rate is to iterate the protocol several times. Specifically, given a protocol Π with FRR of γ , FAR of δ , and complexity C , we can get from it a protocol Π' with smaller FRR as follows: we first run Π once, and if the authentication fails then we run it again. Clearly, Π' has FRR of γ^2 and FAR of 2δ , and we note that its expected complexity is only $C(1 + \gamma)$ (since with probability $1 - \gamma$ we will only run it once).

Similarly, by repeating the original protocol Π upto s times, we obtain a protocol Π^s with FRR of γ^s , FAR of $s\delta$, and for large s the expected complexity will approach $C/(1 - \gamma)$. This means that even protocols with very large false-reject rate (such as the original HB+ that has $FRR \approx 0.44$) are meaningful, in that we can transform them cheaply to protocols with adequately low FRR.

5.5 Protocol Comparison

In Table 2 we compare the trivial exhaustive search with HB+ (ES), our tree-based privacy-preserving HB+ protocol, and the tree-based PRF protocol developed in [14].

Obviously, the calculation on the reader side takes $O(N) \cdot C_{HB+}$ for exhaustive search (where C_{HB+} is the number of calculations for the HB+ authentication protocol), while for the tree-based protocols, this is reduced to logarithmic in N . We also list in that table the storage requirements, the communication, and the corresponding FRR and FAR. (We note that in the exhaustive search case, if the FAR is too high then the FRR

⁶ We note that just like for HB#, it is likely that here too we can reduce the communication to $r(d + 2) + k_x + k_y$ by using Toeplitz matrices for the challenges A, B .

is not meaningful as the probability that a valid tag will be identified and accepted as a different tag is very high.)

For this table we used the values of $\beta = 1000$ for the branching factor of the tree and $d = 2$ for its depth (for a total population of $N = 10^6$ tags). For HB+ we used $r = 80$ for the response length, $\varepsilon = \frac{1}{4}$ for the error rate and $\tau = \varepsilon r = 20$ for the acceptance threshold, and for the size of the keys we used the “low security” values of $k_x = 80$ and $k_y = 256$ (corresponding to the requirement of 2^{55} memory to break LPN). For the PRF we used key size and output size of 128 bits (e.g., for AES-128).

Table 2. Protocol comparison for a population of $N = 10^6$ tags.

Method	Reader Computation	Communication	Tag Memory	FAR	FRR
ES HB+	$10^6 \cdot C_{HB+}$	26960	336	0.98	0.44
Tree HB+	$2000 \cdot C_{HB+}$	27120	848	$4 \cdot 10^{-6}$	0.6
Tree PRF	$2000 \cdot C_{PRF}$	1024	256	0	0

5.6 Choosing the Parameters

With the analysis from above, we now illustrate how to set the parameters of our scheme for a given tag population and security goals. The results of this sections are summarized in Table 3. Specifically, suppose that we want to setup the scheme to achieve the following parameters:

- Tag population of upto $N = 10^6$ tags,
- False-reject rate of 10^{-4} and False-accept rate of 10^{-8} ,
- Security against attacks that work in space of upto 2^{65} bytes.

For these settings, we investigate the parameters that we get by working with noise rates of $\varepsilon = 0.125$ or $\varepsilon = 0.25$, using tree of depth either $d = 2$ or $d = 3$ (and thus $\beta = 1000$ or $\beta = 100$, respectively).

The key-length k_x, k_y . Extrapolating from the parameters in [12, Sec 5.2], for this level of security we need $k_y \approx 440$ for noise level $\varepsilon = 0.125$ and $k_y \approx 330$ for noise level $\varepsilon = 0.25$ (and in either case we can get by with $k_x = 80$).

The response-length r and threshold τ . As mentioned above, we can always get good false-reject rates by repeating the protocol a few times, as long as we start from a protocol that has low enough false-accept rates. Stipulating that repeating the protocol upto four times is reasonable, we thus begin by concentrating on the authentication phase, setting our target for a single run at (say) $\text{FAR} \approx 10^{-9}$ and $\text{FRR} \approx 0.05$.

To get these parameters for the case $\varepsilon = 0.25$ we can use response-length $r = 212$ and threshold $\tau = 63 (= 212/4 + 10)$, which yield $\text{FAR} = 1.62E - 9$ and $\text{FRR} = 0.038$. Similarly, for the $\varepsilon = 0.125$ case we can use $r = 86$ and $\tau = 16 (= 86/8 + 5.25)$, which yield $\text{FAR} = 1.6E - 9$ and $\text{FRR} = 0.036$.

To optimize performance, it is beneficial to set a different response-length for the tree-traversal phase, so as to get $d \cdot \Pr[\text{false branch}] \approx \text{FRR}(\text{auth})$. Below we denote the

response-length for the tree-traversal by r_{tr} . For $\varepsilon = 0.25, \beta = 1000$ we need $r_{\text{tr}} = 102$ (yielding $\Pr[\text{false branch}] = 0.025$), for $\varepsilon = 0.25, \beta = 100$ we need $r_{\text{tr}} = 83$ ($\Pr[\text{false branch}] = 0.0167$), for $\varepsilon = 0.125, \beta = 1000$ we need $r_{\text{tr}} = 40$ ($\Pr[\text{false branch}] = 0.0215$), and for $\varepsilon = 0.25, \beta = 100$ we need $r_{\text{tr}} = 32$ ($\Pr[\text{false branch}] = 0.0146$).

The resulting parameters. The choices above result in four sets of parameters, depending on the values of the noise rate $\varepsilon \in \{0.125, 0.25\}$ and the depth of the tree $d \in \{2, 3\}$ (corresponding to branching factors $\beta \in \{1000, 100\}$, respectively). In either case, running the protocol once induces total communication $r(k_x + k_y) + r_{\text{tr}}d$, computation of $dk_y r_{\text{tr}} + (k_x + k_y)r$ bit operations on the tag and $\beta dk_y r_{\text{tr}} + (k_x + k_y)r$ bit operations on the reader, and storage requirements of $k_x + (d+1)k_y$ for the tags. Running the protocol once with these parameters result in a false-accept rate of roughly 0.1, which is too high for applications, so we repeat it upto four times. As we explained in Section 5.4, this only increases the expected complexity of the protocol by roughly 10%. The result (when running the protocol upto four times) are summarized in Table 3.

Table 3. Concrete parameters for population of $N = 10^6$ tags, running the protocol upto four times to reduce the false-accept rates.

ε	d	β	k_x	k_y	r	r_{tr}	C_{rdr}	C_{tag}	comm	mem	FRR	FAR
0.25	2	1000	80	330	212	102	$7.49E+7$	$1.71E+5$	96804	740	$6.0E-5$	$6.5E-9$
0.25	3	100	80	330	212	83	$9.23E+6$	$1.88E+5$	96854	1400	$6.0E-5$	$6.5E-9$
0.125	2	1000	80	440	86	40	$3.92E+7$	$8.88E+4$	49778	1400	$3.9E-5$	$6.4E-9$
0.125	3	100	80	400	86	32	$4.74E+6$	$9.66E+4$	49796	1840	$4.1E-5$	$6.4E-9$

Legend: ε - error-rate, d - depth, β -branching factor, k_x, k_y - key lengths, r - response length in auth. stage, r_{tr} - response length in tree stage, C_{rdr} - expected reader computation, C_{tag} - expected tag computation, **comm** - expected total communication, **mem** - tag memory requirements.

It can be seen from Table 3 that most of the relevant parameters are improved significantly when moving to a lower-error regime (i.e., from $\varepsilon = 0.25$ to $\varepsilon = 0.125$), and also when moving to deeper trees (i.e., from $d = 2$ to $d = 3$). The only limiting parameter is the memory requirement at the tag, which increases for deeper trees and smaller error rates. Therefore, for any particular application, one should use as deep a tree and as small an error as can be realized subject to the memory available at the tags.

We again note that the communication complexity of the protocol (which is rather large) can probably be significantly reduced by using Toeplitz matrices for the challenges A, B .⁷ (In the examples above, this optimization will reduce the communication complexity from the current range of 50000-100000 bits to only about 1000 bits.) Another “big win” will be to be able derive the keys \mathbf{y} pseudo-randomly manner from shorter secrets (e.g., using a cheap PRG such as the shrinking generator [5]).

⁷ Another alternative is to replace $A\mathbf{x} + B\mathbf{y}$ with $\mathbf{a} \cdot \mathbf{x} + \mathbf{b} \cdot \mathbf{y}$ over a finite field.

6 Simulation

To estimate the computational overhead incurred at the reader-side when using different methods, we performed a simulation. Our simulation was done using Matlab. The computer used for the simulation is an IBM thinkpad, with Intel CPU, T2400, 1.83 GHz and 1.99 GB RAM.

The test was run for a million tags. We tested two cases for our Tree-HB+ protocol: A two-level tree with 1000 nodes in each level, and a three-level tree with 100 nodes at each level, all with the parameters that were proposed in the original HB+ protocol (i.e., error-rate of $\varepsilon = 0.25$, using $k_x = k_y = 224$). Our results show that for $r = 96$, the exhaustive search average run-time using HB+ was 110 seconds. For the tree-based search (2 levels), we get an average run time of 0.1307 seconds, and for 3 levels we get run time of 0.019. This ratio is compatible with our theoretical estimates, as the exhaustive search will take $O(N)$ (e.g., $\approx 10^6$) while the tree based search only takes $O(d \log N)$ (e.g., $\approx 2 \cdot 10^3$ or $\approx 3 \cdot 10^2$, respectively). The AES tree-based search was performed with a 16-byte random key, and is using the Java Crypto libraries to perform the actual AES calculations. The run time for the tree-based search using AES was around 7 seconds, which confirms the fact that AES requires significantly more calculations than standard HB+ protocol. We can further estimate that since the Meet-in-the-Middle PRF scheme takes about $\frac{2}{\sqrt{\beta}} = 0.063$ of the tree-based PRF scheme, it would take about 0.44 seconds to run it on our computer. Therefore, our simulation shows that our privacy-preserving method is faster than the Tree-based PRF and the Meet-in-the-Middle strategy. For $d = 3$, our method is faster by a factor of ≈ 53 relative to the Tree-based PRF and ≈ 3 comparing to the MITM method (for $d = 2$, we get the factors of roughly 370 and 23, respectively).

Table 4. Simulation Results, $r = 96$

Method	Comp. Time (seconds) mean \pm stdev	FAR expected/observed	FRR expected/observed
ES (HB+)	110.180 \pm 0.9656	0.76 / 0.75	0.365 / 0.450
Tree HB+ ($d = 2$)	0.131 \pm 0.0234	1.4E-6 / 0	0.432 / 0.422
Tree HB+ ($d = 3$)	0.019 \pm 6.75E-4	1.4E-6 / 2.14E-6	0.382 / 0.382
Tree PRF(AES)	7.046 \pm 0.0561	—	—

The results of our simulation are shown in Table 5. The mean and standard deviation of the run-time are provided, together with the FAR and FRR for each case. For the privacy-preserving Tree HB+ protocol, the FAR results were not accurate due to the fact that the expected FAR is only $1.4 \cdot 10^{-6}$ and therefore a very large number of runs are needed to get accurate results.

7 Conclusions and Future Work

In this paper we developed tree-based HB-type protocols for privacy-preserving authentication. This is useful for two reasons. First, very low-cost tags may be incapable of

computing standard PRFs and HB-type protocols are currently the only viable solution for such tags. Second, since the underlying computations in HB protocols are very efficient, it automatically reduces reader load compared to PRF-based protocols.

We proposed some significant improvements over the naive use of HB+ in a tree-based scheme. These improvements reduce the computation and communication by a factor of $2\times$ to $4\times$. In fact our tree protocol is almost as efficient as the underlying HB+. The error rates in our protocols are nearly as low as that underlying HB+ protocol. This makes our scheme suitable for a system with large number of tags (Unlike exhaustive search, which would produce unacceptable security for practical systems). We presented analytical and simulation results comparing our method with prior proposals in terms of computation, communication and memory overheads.

References

1. B. Applebaum. Fast Cryptographic Primitives Based on the Hardness of Decoding Random Linear Code. Princeton University Technical Report TR-845-08, December 2008.
2. M. Burmester, B. Medeiros and R. Motta. Robust anonymous RFID authentication with constant key-lookup. ACM symposium on Information, Computer and Communications Security (ASIACCS '08). pp. 283-291. 2008.
3. A. Blum, A. Kalai and H. Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. J. ACM 50(4): 506-519, 2003.
4. J.H. Cheon, J. Hong and G. Tsudik. Reducing RFID Reader Load with the meet-in-the-Middle Strategy, IACR ePrint report 2009/271, 2009.
5. D. Coppersmith, H. Krawczyk and Y. Mansour. The Shrinking Generator. CRYPTO '93, LNCS vol. 773, pages 22-39. Springer, 1993.
6. H. Gilbert, M.J.B. Robshaw, and Y. Seurin. HB#: Increasing the Security and Efficiency of HB+, Eurocrypt '08. LNCS vol. 4965, pages 361-378. Springer, 2008.
7. H. Gilbert, M.J.B. Robshaw, and H. Sibert. Active attack against HB+: a provably secure lightweight authentication protocol. IEE Electronic Letters 41(21), pages 1169-1170, 2005.
8. N.J. Hopper and M. Bulm. Secure human identification protocols. Asiacrypt '01, LNCS vol. 2248, pages 52-66. Springer, 2001.
9. A. Juels and S. Weis. Authenticating Pervasive Devices with Human Protocols CRYPTO '05, LNCS vol. 3621, pages 293-308. Springer, 2005.
10. J. Katz and J.S. Shin. Parallel and Concurrent Security of the HB and HB+ Protocols. EUROCRYPT '06, LNCS vol. 4004, pages 73-86. Springer, 2006.
11. J. Katz, A. Smith. Analyzing the HB and HB+ Protocols in the "Large Error" case. IACR ePrint report 2006/326, 2006.
12. E. Levieil and P.A. Fouque. An Improved LPN algorithm. SCN '06, LNCS vol. 4119, pages 348-359. Springer, 2006.
13. D. Molnar, A. Soppera and D. Wagner. A Scalable Delegetable Pseudonym Protocol Enable Ownership Transfer of RFID Tags. The 2nd ACM conference on Wireless network security, 2009.
14. D. Molnar and D. Wagner. Privacy and security in library RFID: issues, practices ACM-CCS '04, pages 210-219, ACM 2004.
15. Ben Morris, Phillip Rogaway, Till Stegers. How to Encipher Messages on a Small Domain: Deterministic Encryption and the Thorp Shuffle. CRYPTO '09, LNCS vol. 5677, pages 286-302. Springer, 2009.
16. K. Quafi, R. Overbock. S. Vaudenay. On the Security of HB# against a Man-in-the-Middle Attack. Asiacrypt '08, LNCS vol. 5350, pages 108-124. Springer, 2008.